On Conflict-Free Chromatic Guarding of Simple Polygons

Onur Çağırıcı[1], Subir Kumar Ghosh[2], Petr Hliněný[1](✉), and Bodhayan Roy[3]

[1] Faculty of Informatics of Masaryk University, Brno, Czech Republic
onur@mail.muni.cz, hlineny@fi.muni.cz

[2] Ramakrishna Mission Vivekananda University, Kolkata, India
ghosh@tifr.res.in

[3] Indian Institute of Technology Kharagpur, Kharagpur, India
broy@maths.iitkgp.ac.in

Abstract. We study the problem of colouring the vertices of a polygon, such that every viewer can see a unique colour. The goal is to minimize the number of colours used. This is also known as the conflict-free chromatic guarding problem with vertex guards, and is motivated, e.g., by the problem of radio frequency assignment to sensors placed at the polygon vertices. We study the scenario in which viewers can be all points of the polygon (such as a mobile robot which moves in the interior of the polygon). We efficiently solve the related problem of minimizing the number of guards and approximate (up to only an additive error) the number of colours required in the special case of polygons called funnels. As a corollary we sketch an upper bound of $O(\log^2 n)$ colours on n-vertex weak visibility polygons which generalizes to all simple polygons.

Keywords: Computational geometry · Polygon guarding · Visibility graph · Art gallery problem · Conflict-free coloring

1 Introduction

The *guarding* of a polygon is placing "guards" into the polygon, in a way that the guards collectively can see the whole polygon. It is usually assumed that a guard can see any point unless there is an obstacle or a wall between the guard and that point. One of the best known problems in computational geometry, the *art gallery problem* is essentially a guarding problem [7,28]. The problem is to find the minimum number of guards to guard an art gallery, which is

O. Çağırıcı—Supported by the Czech Science Foundation, project no. 17-00837S.
S. K. Ghosh—Supported by SERB, Government of India through a grant under MATRICS.
P. Hliněný—Supported by the Czech Science Foundation, project no. 17-00837S.
B. Roy—A significant part of the work was done while the author was affiliated to the Faculty of Informatics of Masaryk University.

Y. Li et al. (Eds.): COCOA 2019, LNCS 11949, pp. 601–612, 2019.
https://doi.org/10.1007/978-3-030-36412-0_49

modeled by an n-vertex polygon. This problem was shown to be NP-hard by Lee and Lin [27] and more recently $\exists\mathbb{R}$-complete by Abrahamsen et al. [2]. The Art Gallery Theorem, proved by Chvátal, shows that $\lfloor n/3 \rfloor$ guards are sufficient and sometimes necessary to guard a simple polygon [15].

The guard minimization problem has been studied under many constraints; such as the placement of guards being restricted to the polygonal perimeter or vertices [25], the viewers being restricted to vertices, the polygon being terrains [3,6,16], weakly visible from an edge [8], with holes or orthogonal [10,17,24], with respect to parameterization [11], approximability [23].

For most of these cases the problem remains hard, but interesting approximation algorithms have also been provided [12,19].

In addition to above mentioned versions of art gallery problem (or rather polygon guarding problem), some problems consider not the number of the guards, but the number of colours that are assigned to the guards. The colours, depending on the scope, determine the types of the guards. If any observer in the polygon sees at least one guard with a different type, then that polygon has a *conflict-free chromatic guarding* [4,5,22]. If every guard that any given observer sees is of different type, then that polygon has a *strong chromatic guarding* [18].

Motivation. In general, conflict-free colouring of a graph is assigning colours to vertices of that graph such that the neighborhood of each vertex contains at least one unique colour. This problem was first studied by Biggs with the name *perfect code*, which is essentially conflict-free colouring of a graph using only one colour [9,26]. Later on, this topic aroused interest on polygon visibility graphs when the field of robotics became widespread [14,20].

Consider a scenario where a mobile robot traverses a room from one point to another, communicating with the wireless sensors placed on the corners of the room. Even if the robot has full access to the map of the room, it cannot determine its location precisely because of accumulating rounding errors. And thus it needs clear markings in the room to guide itself to the end point in an energy efficient way. To guide a mobile robot with wireless sensors, two properties must be satisfied. First one is, no matter where the robot is in the polygon, it should hear from at least one sensor. That is, the placed sensors must together *guard* the whole room and leave no place uncovered. The second one is, if the robot hears from several sensors, there must be at least one sensor broadcasting with a frequency that is not reused by some other sensor in the range. That is, the sensors must have *conflict-free* frequencies. If these two properties are satisfied, then the robot can guide itself using the deployed wireless sensors as landmarks. This problem is also closely related to frequency assignment problem in wireless networks [1,5]. One can easily solve this problem by placing a sensor at each corner of the room, and assigning a different frequency to each sensor. However, this method becomes very expensive as the number of sensors grow [1,29]. Therefore, the main goal in this problem is minimize the number of different frequencies assigned to sensors. Since the cost of a sensor is comparatively very low, we do not aim to minimize the number of sensors used.

The above scenario is geometrically modeled as follows. The room is a simple polygon with n vertices. There are m sensors placed in the polygon (usually on some of its vertices), and two different sensors are given two different colours if, and only if they broadcast in different frequencies.

Basic Definitions. We consider simple polygons (informally, "without holes"), usually non-convex. Two points p_1 and p_2 of a polygon P are said to *see* each other, or be *visible* to each other, if the line segment $\overline{p_1 p_2}$ fully belongs to P. In this context, we say that a guard g *guards* a point x of P if the line segment $\overline{gx}$ fully belongs to P. A polygon P is a *weak visibility polygon* if P has an edge uv such that for every point p of P there is a point p' on uv seeing p.

In the paper, we pay close attention to a special type of polygons – *funnels*. A polygon P is a funnel if (Fig. 1) precisely three of the vertices of P are convex, and two of the convex vertices share one common edge – the *base* of the funnel. Funnels attract special interest in the study of visibility graphs, as a very fundamental class of polygons. Other polygons can be decomposed into funnels, giving a good overview of their structure with respect to most geometric problems. Funnels have a simpler structure due to their two concave chains and hence allow for easier handling of visibility problems than other classes of polygons.

A solution of *conflict-free chromatic guarding* of a polygon P consists of a set of *guards* in P, and an assignment of colours to the guards (one colour per guard) such that the following holds; every *viewer* v in P (where v can be any point of P in our case) can see a guard of colour c such that no other guard seen by v has the same colour c. In the *point-to-point* (P2P) variant the guards can be placed in any points of P, while in the *vertex-to-point* (V2P) variant the guards can be placed only at the vertices of P. (There also exists a V2V variant in which also viewers are restricted to the vertices.) In all variants the goal is to minimize the number of colours (e.g., frequencies) used.

When writing $\log n$, we mean the binary logarithm $\log_2 n$.

Related Research. The aforementioned P2P conflict-free chromatic guarding (art gallery) problem has been studied in several papers. Bärtschi and Suri gave an upper bound of $O(\log^2 n)$ colours on simple n-vertex polygons [5]. Later, Bärtschi et al. improved this upper bound to $O(\log n)$ on simple polygons [4], and Hoffmann et al. [22], while studying the orthogonal variant of the problem, have given the first nontrivial lower bound of $\Omega\big(\log\log n/\log\log\log n\big)$ colours holding also in the general case of simple polygons.

Our paper deals with the V2P variant in which guards should be placed on polygon vertices and viewers can be any points of the polygon. Note that there are some fundamental differences between point and vertex guards, e.g., funnel polygons (Fig. 1) can always be guarded by one point guard (of one colour) but they may require up to $\Omega(\log n)$ colours in the V2P conflict-free chromatic guarding, as shown in [4]. Hence, extending a general upper bound of $O(\log n)$ colours for point guards on simple polygons by Bärtschi et al. [4] to the more restrictive vertex guards is a challenge, which we can now only approach with

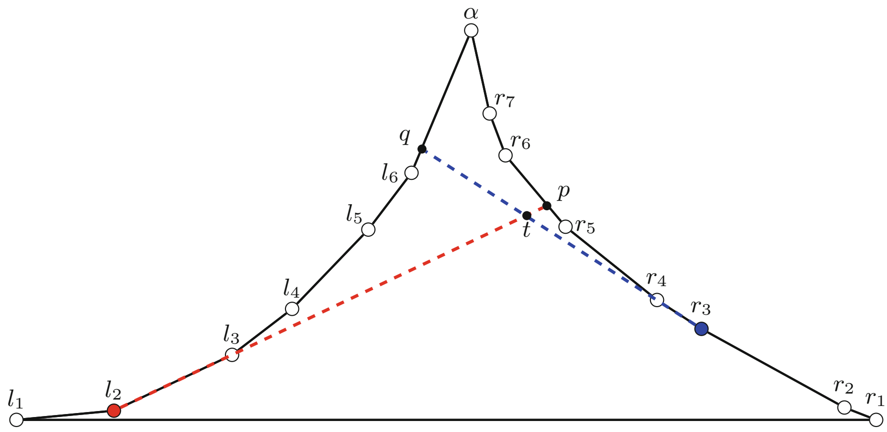

Fig. 1. A funnel F with seven vertices in $\mathcal{L}$ labeled $l_1, \ldots, l_7$ from bottom to top, and eight vertices in $\mathcal{R}$ labeled $r_1, \ldots, r_8$, including the apex $\alpha = l_7 = r_8$. The picture also shows the upper tangent of the vertex l_2 of $\mathcal{L}$ (drawn in dashed red), the upper tangent of the vertex r_3 of $\mathcal{R}$ (drawn in dashed blue), and their intersection t. (Color figure online)

an $O(\log^2 n)$ bound (see below). Note that the same bound of $O(\log^2 n)$ colours was attained by [4] when allowing multiple guards at the same vertex.

Our Results. We give a polynomial-time algorithm to find the optimum number m of vertex-guards to guard all the points of a funnel, and show that the number of colours in the corresponding conflict-free chromatic guarding problem is $\log m + \Theta(1)$ (Theorem 2). This leads to an approximation algorithm for V2P conflict-free chromatic guarding of a funnel, with only a constant (+4) additive error (Corollary 9). A remarkable feature of this result is that we prove a direct relation (Theorem 8) between the optimal numbers of guards and of colours needed in funnel polygons. Finally, we sketch that a weak visibility polygon on n vertices can be V2P conflict-free chromatic guarded with only $O(\log^2 n)$ guards, and generalize this upper bound to all simple polygons, which is a result incomparable with previous [4].

Note that all our algorithms are simple and suitable for easy implementation. Due to space restrictions, the proofs of a part of the statements are left for the full paper published on arXiv [13], and those statements are marked with (*).

2 Minimizing Vertex-to-Point Guards for Funnels

Before turning to the conflict-free chromatic guarding problem, we first resolve the problem of minimizing the total number of vertex guards needed to guard all points of a funnel polygon. We start by describing a simple procedure (Algorithm 1) that provides us with a guard set which may not always be optimal

Algorithm 1. Simple vertex-to-point guarding of funnels (uncoloured).

Input: A funnel F with concave chains $\mathcal{L} = (l_1, \ldots, l_k)$ and $\mathcal{R} = (r_1, \ldots, r_m)$.
Output: A vertex set guarding all the points of F.

Initialize an auxiliary digraph G with two dummy vertices x and y, and declare $ups(x) = \overline{l_1 r_1}$;
Initialize $\mathcal{S} \leftarrow \{x\}$;
while $\mathcal{S}$ *is not empty* **do**
  Choose an arbitrary $t \in \mathcal{S}$, and remove t from $\mathcal{S}$;
  Let $s = ups(t)$; /* s is a segment inside F */
  Let q and p be the ends of s on $\mathcal{L}$ and $\mathcal{R}$, respectively;
  Let i and j be the largest indices such that l_i and r_j are not above q and p, resp.;
  if l_{i+1} *can see whole* s **then** $i' \leftarrow i + 1$;
  else $i' \leftarrow i$; /* the topmost vertex on the left seeing whole s */
  if r_{j+1} *can see whole* s **then** $j' \leftarrow j + 1$;
  else $j' \leftarrow j$; /* the topmost vertex on the right seeing whole s */
  Include the vertices $l_{i'}$ and $r_{j'}$ in G;
  foreach $z \in \{l_{i'}, r_{j'}\}$ **do**
    Add the directed edge (t, z) to G ;
    if *segment* $ups(z)$ *includes the apex* $l_k = r_m$ **then**
      Add the directed edge (z, y) to G ; /* y is the dummy vertex */
    else $\mathcal{S} \leftarrow \mathcal{S} \cup \{z\}$; /* more guards are needed above z */
Enumerate a shortest path from x to y in G;
Output the shortest path vertices without x and y as the required guard set;

(but very close to the optimum, see Corollary 3). This procedure will be helpful for the subsequent colouring results. Then we also refine the simple procedure to compute the optimal number of guards in Algorithm 2.

We use some special notation here. See Fig. 1. Let the given funnel be F, oriented in the plane as follows. On the bottom, there is the horizontal *base* of the funnel – the line segment $\overline{l_1 r_1}$ in the picture. The topmost vertex of F is called the *apex*, and it is denoted by α. There always exists a point x on the base which can see the apex α, and then x sees the whole funnel at once. The vertices on the left side of apex form the *left concave chain*, and analogously, the vertices on the right side of the apex form the *right concave chain* of the funnel. These left and right concave chains are denoted by $\mathcal{L}$ and $\mathcal{R}$ respectively. We denote the vertices of $\mathcal{L}$ as $l_1, l_2, \ldots, l_k$ from bottom to top. We denote the vertices of $\mathcal{R}$ as $r_1, r_2, \ldots, r_m$ from bottom to top. Hence, the apex is $l_k = r_m = \alpha$.

Let l_i be a vertex on $\mathcal{L}$ which is not the apex. We define the *upper tangent* of l_i, denoted by $upt(l_i)$, as the ray whose origin is l_i and which passes through l_{i+1}. Upper tangents for vertices on $\mathcal{R}$ are defined analogously. Let p be the point of intersection of $\mathcal{R}$ and the upper tangent of l_i. Then we define $ups(l_i)$ as the line segment $\overline{l_{i+1} p}$. For the vertices of $\mathcal{R}$, ups is defined analogously: if q is the point of intersection of $\mathcal{L}$ and the upper tangent of $r_j \in \mathcal{R}$, then let $ups(r_j) := \overline{r_{j+1} q}$.

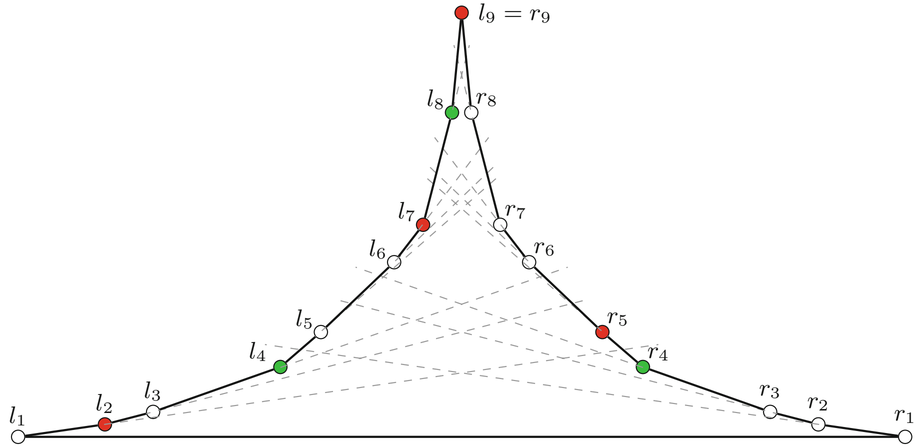

Fig. 2. A symmetric funnel with 17 vertices. The gray dashed lines show the upper tangents of the vertices. It is easy to see that Algorithm 1 selects 4 guards, up to symmetry, at l_2, r_5, l_7, l_9 (the red vertices). However, the whole funnel can be guarded by three guards at l_4, r_4, l_8 (the green vertices). (Color figure online)

The underlying idea of Algorithm 1 is as follows. Imagine we proceed bottom-up when building the guard set of a funnel F. Then the next guard is placed at the top-most vertex z of F, nondeterministically choosing between z on the left and the right chain of F, such that no "unguarded gap" remains below z. Note that the unguarded region of F after placing a guard at z is bounded from below by $ups(z)$. The nondeterministic choice of the next guard z is encoded within a digraph, in which we then find the desired guard set as a shortest path.

Lemma 1 (*). *Algorithm 1 runs in polynomial time, and it outputs a feasible guard set for all the points of a funnel F.*

Unfortunately, the guard set produced by Algorithm 1 may not be optimal under certain circumstances. See the example in Fig. 2; the algorithm picks the four red vertices, but the funnel can be guarded by the three green vertices.

For the sake of completeness, we now refine the simple approach of Algorithm 1 to always produce a minimum size guard set. Our refinement is going to consider also pairs of guards (one from the left and one from the right chain) in the procedure. We correspondingly extend the definition of *ups* to pairs of vertices as follows. Let l_i and r_j be vertices of F on $\mathcal{L}$ and $\mathcal{R}$, respectively, such that $ups(l_i) = \overline{l_{i+1}p}$ intersects $ups(r_j) = \overline{r_{j+1}q}$ in a point t (see in Fig. 1). Then we set $ups(l_i, r_j)$ as the polygonal line ("$\vee$-shape") $\overline{pt} \cup \overline{qt}$.

Algorithm 2, informally saying, enriches the two nondeterministic choices of placing the next guard in Algorithm 1 with a third choice; placing a suitable top-most pair of guards $z = (z_1, z_2)$, $z_1 \in \mathcal{L}$ and $z_2 \in \mathcal{R}$, such that again no "unguarded gap" remains below (z_1, z_2). Figure 2 features a funnel in which placing such a pair of guards ($z_1 = l_4$, $z_2 = r_4$) may be strictly better than using

Algorithm 2. (*) Optimum vertex-to-point guarding of funnels.

Input: A funnel F with concave chains $\mathcal{L} = (l_1, \ldots, l_k)$ and $\mathcal{R} = (r_1, \ldots, r_m)$.
Output: A minimum vertex set (uncoloured) guarding all the points of F.

* On line 13 of Algorithm 1, consider $z \in \{l_{i'}, r_{j'}, (l_{i''}, r_{j''})\}$, where i'' and j'' are the largest indices such that $l_{i''}$ lies strictly below $ups(p)$ and $r_{j''}$ strictly below $ups(q)$. (Then $l_{i''}$ and $r_{j''}$ together can see whole s.);
* On line 14 of Algorithm 1, make the edge (t, z) of G weight 2 if $z = (l_{i''}, r_{j''})$.

any two consecutive steps of Algorithm 1. On the other hand, we can show that there is no better possibility than one of these three considered steps, giving us:

Theorem 2 (*). *Algorithm 2 runs in polynomial time, and it outputs a feasible guard set of minimum size guarding all the points of a funnel F.*

Lastly, we establish that the difference between Algorithms 1 and 2 cannot be larger than 1 guard. Let G^1 with the source x^1 be the auxiliary graph produced by Algorithm 1, and G^2 with the source x^2 be the one produced by Algorithm 2. We can prove the following detailed statement by induction on $i \geq 0$:

- Let $P^2 = (x^2 = x_0^2, x_1^2, \ldots, x_i^2)$ be any directed path in G^2 of weight k, let Q^2 denote the set of guards listed in the vertices of P^2, and $L^2 = \mathcal{L} \cap Q^2$ and $R^2 = \mathcal{R} \cap Q^2$. Then there exists a directed path $(x^1 = x_0^1, x_1^1, \ldots, x_k^1, x_{k+1}^1)$ in G^1 (of length $k+1$), such that the guard of x_k is at least as high as all the guards of L^2 (if $x_k \in \mathcal{L}$) or of R^2 (if $x_k \in \mathcal{R}$), and the guard of x_{k+1} is strictly higher than all the guards of Q^2.

Corollary 3 (*). *The guard set produced by Algorithm 1 is always by at most one guard larger than the optimum solution produced by Algorithm 2.*

3 V2P Conflict-Free Chromatic Guarding of Funnels

In this section, we continue to study funnels. To obtain a conflict-free coloured solution, we will simply consider the guards chosen by Algorithm 1 in the ascending order of their vertical coordinates, and colour them in the *ruler sequence*, (e.g., [21]) in which the i^{th} term is the exponent of the largest power of 2 that divides $2i$. (The first few terms are $1, 2, 1, 3, 1, 2, 1, 4, 1, 2, 1, 3, 1, 2, 1, 5, 1 \ldots$.) So, if Algorithm 1 gives m guards, then our approach will use about $\log m$ colours.

Our aim is to show that this is always very close to the optimum, by giving a lower bound on the number of necessary colours of order $\log m - O(1)$. To achieve this, we study the following two sets of guards for a given funnel F:

- The minimal *guard set* A computed by Algorithm 1 on F (which is overall nearly optimal by Corollary 3); if this is not unique, then we fix any such A.
- A *guard set* D which achieves the minimum number of colours for conflict-free guarding; note that D may be much larger than A since it is the number of colours which matters.

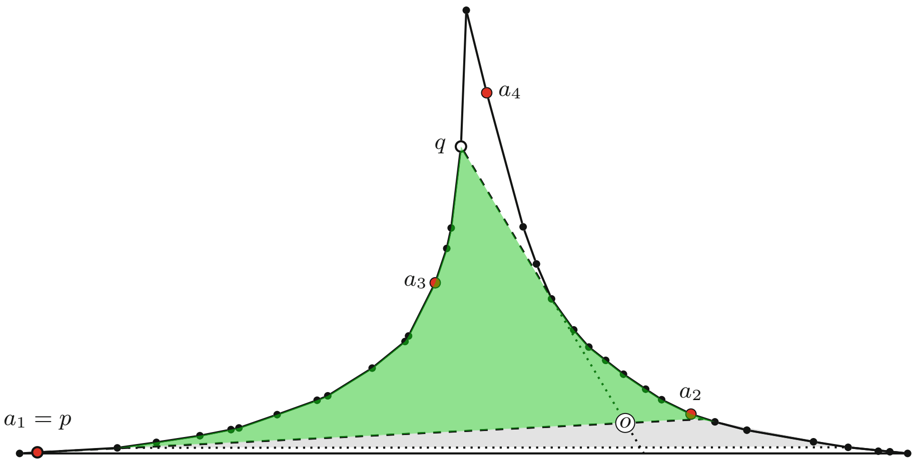

Fig. 3. An example of a 2-interval Q of a funnel (green and bounded by $s_1 = ups(p)$ and $s_2 = los(q)$). The red vertices $a_1 = p, a_2, a_3, a_4$ are the guards computed by Algorithm 1, and a_2, a_3 belong to the interval Q. The shadow of Q (filled light gray) is bounded from below by the bottom dotted line, and the inner point o is the so-called *observer* of Q. (Color figure online)

On a high level, we are going to show that the colouring of D must (somehow) copy the ruler sequence on A. For that we will recursively bisect our funnel into smaller "layers", gaining one unique colour with each bisection.

Analogously to the notion of an upper tangent from Sect. 2, we define the *lower tangent* of a vertex $l_i \in \mathcal{L}$, denote by $lot(l_i)$, as the ray whose origin is l_i and which passes through $r_j \in \mathcal{R}$ such that r_j is the lowest vertex on $\mathcal{R}$ seeing l_i. Note that $lot(l_i)$ may intersect $\mathcal{R}$ in r_j alone or in a segment from r_j up. Let $los(l_i) := \overline{l_i r_j}$. The definition of $lot()$ and $los()$ for vertices of $\mathcal{R}$ is symmetric.

We now give a definition of "layers" of a funnel which is crucial for our proof.

Definition 4 (t-interval). Let F be a funnel with the chains $\mathcal{L} = (l_1, l_2, \ldots, l_k)$ and $\mathcal{R} = (r_1, r_2, \ldots, r_m)$, and A be the fixed guard set A computed by Algorithm 1 on F. Let s_1 be the base of F, or $s_1 = ups(p)$ for some vertex p of F (where p is not the apex or its neighbour). Let s_2 be the apex of F, or $s_2 = los(q)$ for some vertex q of F (where q is not in the base of F). Assume that s_2 is above s_1 within F. Then the region Q of F bounded from below by s_1 and from above by s_2, excluding q itself, is called an *interval of* F. Moreover, Q is called a *t-interval of* F if Q contains at least t of the guards of A. See Fig. 3.

Having an interval Q of the funnel F, bounded from below by s_1, we define the *shadow of* Q as follows. If $s_1 = ups(l_i)$ ($s_1 = ups(r_j)$), then the shadow consists of the region of F between s_1 and $los(l_{i+1})$ (between s_1 and $los(r_{j+1})$, respectively). If s_1 is the base, then the shadow is empty.

Lemma 5 (*). *If Q is a 13-interval of the funnel F, then there exists a point in Q which is not visible from any vertex of F outside of Q.*

Our second crucial ingredient is the possibility to "almost privately" see the vertices of an interval Q from one point as follows. If $s_2 = los(q)$, then the intersection point of $lot(q)$ with s_1 is called the *observer of* Q. (Actually, to be precise, we should slightly perturb this position of the observer o so that the visibility between o and q is blocked.) If s_2 is the apex, then consider the spine of F instead of $lot(q)$. See again Fig. 3. The following is easy to argue.

Lemma 6 (*). *The observer o of an interval Q in a funnel F can see all the vertices of Q, but o cannot see any vertex of F which is not in Q and not in the shadow of Q.*

The last ingredient before the main proof is the notion of sections of an interval Q of F. Let s_1 and s_2 form the lower and upper boundary of Q. Consider a vertex $l_i \in \mathcal{L}$ of Q. Then the *lower section of Q at l_i* is the interval of F bounded from below by s_1 and from above by $los(l_i)$. The *upper section of Q at l_i* is the interval of F bounded from below by $ups(l_i)$ and from above by s_2. Sections of $r_j \in \mathcal{R}$ are defined analogously. Again, the following is straightforward.

Lemma 7 (*). *Let Q be a t-interval of the funnel F, and let Q_1 and Q_2 be its lower and upper sections at some vertex p. Then Q_i, $i = 1, 2$, is a t_i-interval such that $t_1 + t_2 \geq t - 3$.*

Theorem 8. *Any conflict-free chromatic guarding of a given funnel requires at least $\lfloor \log_2(m+3) \rfloor - 3$ colours, where m is the minimum number of guards needed to guard the whole funnel.*

Proof. We will prove the following claim by induction on $c \geq 0$: If Q is a t-interval in the funnel F and $t \geq 16 \cdot 2^c - 3$, then any conflict-free colouring of F must use at least $c + 1$ colours on the vertices of Q or of the shadow of Q.

In the base $c = 0$ of the induction, we have $t \geq 16 - 3 = 13$. By Lemma 5, some point of Q is not seen from outside, and so there has to be a coloured guard in some vertex of Q, thus giving $c + 1 = 1$ colour.

Consider now $c > 0$. The observer o of Q (which sees all the vertices of Q) must see a guard g of a unique colour where g is, by Lemma 6, a vertex of Q or of the shadow of Q. In the first case, we consider Q_1 and Q_2, the lower and upper sections of Q at g. By Lemma 7, for some $i \in \{1, 2\}$, Q_i is a t_i-interval of F such that $t_i \geq (t-3)/2 \geq (16 \cdot 2^c - 6)/2 = 16 \cdot 2^{c-1} - 3$. In the second case ($g$ is in the shadow of Q), we choose g' as the lowermost vertex of Q on the same chain as g, and take only the upper section Q_1 of Q at g'. We continue as in the first case with $i = 1$.

By induction assumption for $c-1$, Q_i together with its shadow carry a set C of at least c colours. The shadow of Q_2 is included in Q, and the shadow of Q_1 coincides with the shadow of Q, moreover, the observer of Q_1 sees only a subset of the shadow of Q seen by the observer o of Q. Since g is not a point of Q_i or its shadow, but our observer o sees the colour c_g of g and all the colours of C, we have $c_g \notin C$ and hence $C \cup \{c_g\}$ has at least $c + 1$ colours, as desired.

Finally, we apply the above claim to $Q = F$. We have $t \geq m$, and for $t \geq m \geq 16 \cdot 2^c - 3$ we derive that we need at least $c + 1 \geq \lfloor \log(m + 3) \rfloor - 3$ colours for guarding whole F. □

Algorithm 3. Approximate conflict-free chromatic guarding of a funnel.

Input: A funnel F with concave chains $\mathcal{L} = (l_1, \ldots, l_k)$ and $\mathcal{R} = (r_1, \ldots, r_m)$.
Output: A conflict-free chromatic guard set of F using $\leq OPT + 4$ colours.

1 Run Algorithm 1 to produce a guard seq. $A = (a_1, a_2, \ldots, a_t)$ (bottom-up);
2 Assign colours to members of A according to the ruler sequence; the vertex a_i gets colour c_i where c_i is the largest integer such that 2^{c_i} divides $2i$;
3 Output coloured guards A as the (approximate) solution;

Corollary 9. *Algorithm 3, for a given funnel F, outputs in polynomial time a conflict-free chromatic guard set A, such that the number of colours used by A is by at most four larger than the optimum.*

Proof. Note the following simple property of the ruler sequence: if $c_i = c_j$ for some $i \neq j$, then $c_{(i+j)/2} > c_i$. Hence, for any i, j, the largest value occurring among colours $c_i, c_{i+1}, \ldots, c_{i+j-1}$ is unique. Since every point of F sees a consecutive subsequence of A, this is a feasible conflict-free colouring of the funnel F.

Let m be the minimum number of guards in F. By Corollary 3, it is $m+1 \geq t = |A| \geq m$. To prove the approximation guarantee, observe that for $t \leq 2^c - 1$, our sequence A uses $\leq c$ colours. Conversely, if $t \geq 2^{c-1}$, i.e. $m \geq 2^{c-1} - 1$, then the required number of colours for guarding F is at least $c - 1 - 3 = c - 4$, and hence our algorithm uses at most 4 more colours than the optimum. □

4 Concluding Remarks

We have designed an algorithm for producing a V2P guarding of funnels that is optimal in the number of guards. We have also designed an algorithm for a V2P conflict-free chromatic guarding for funnels, which gives only an additive error (+4) with respect to the minimum number of colours required. We believe that the latter can be strengthened to an exact solution by sharpening the arguments involved (though, it would likely not be easy).

Regarding V2P conflict-free chromatic guarding in a more general setting, we provide the following upper bound as a corollary of the previous result.

Theorem 10 (*). *There is an algorithm computing in polynomial time a conflict-free chromatic guarding of a weak visibility polygon using $O(\log^2 n)$ colours.*

A rough sketch of the proof is as follows. Each weak visibility polygon can be straightforwardly partitioned into a sequence of maximal (overlapping) funnels, and we can independently guard each of the funnels with $O(\log n)$ colours by applying Algorithm 3. These colourings, unsurprisingly, may conflict with each other, and so we additionally couple the colours in funnels with $O(\log n)$ colours of the ruler sequence assigned to each of the funnels.

Secondly, we can use a polygon decomposition technique introduced by Suri [30] to generalize the upper bound from weak visibility polygons (Theorem 10) to all simple polygons. The technique has already been used by Bärtschi et al. in [4] for the P2P version of our problem. Though, there is still a room for improvement down to $O(\log n)$, which is the worst-case scenario already for funnels and which would match the previous P2P upper bound for simple polygons of [4].

To summarize, we propose the following open problems for future research:

– Improve Corollary 9 to an exact algorithm for guarding a funnel.
– Improve the upper bound in Theorem 10 to $O(\log n)$.

References

1. Aardal, K.I., van Hoesel, S.P.M., Koster, A.M.C.A., Mannino, C., Sassano, A.: Models and solution techniques for frequency assignment problems. Ann. Oper. Res. **153**, 79–129 (2007)
2. Abrahamsen, M., Adamaszek, A., Miltzow, T.: The art gallery problem is $\exists \mathbb{R}$-complete. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, pp. 65–73 (2018)
3. Ashok, P., Fomin, F.V., Kolay, S., Saurabh, S., Zehavi, M.: Exact algorithms for terrain guarding. ACM Trans. Algorithms **14**, 1–20 (2018)
4. Bärtschi, A., Ghosh, S.K., Mihalák, M., Tschager, T., Widmayer, P.: Improved bounds for the conflict-free chromatic art gallery problem. In: Proceedings of the 30th Annual Symposium on Computational Geometry, pp. 144–153 (2014)
5. Bärtschi, A., Suri, S.: Conflict-free chromatic art gallery coverage. Algorithmica **68**, 265–283 (2014)
6. Ben-Moshe, B., Katz, M.J., Mitchell, J.S.B.: A constant-factor approximation algorithm for optimal 1.5D terrain guarding. SIAM J. Comput. **36**, 1631–1647 (2007)
7. de Berg, M., Cheong, O., Kreveld, M., Overmars, M.: Computational Geometry, Algorithms and Applications, 3rd edn. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-77974-2
8. Bhattacharya, P., Ghosh, S.K., Roy, B.: Approximability of guarding weak visibility polygons. Discrete Appl. Math. **228**, 109–129 (2017)
9. Biggs, N.: Perfect codes in graphs. J. Comb. Theory Ser. B **15**, 289–296 (1973)
10. Bonnet, É., Giannopoulos, P.: Orthogonal terrain guarding is NP-complete. In: 34th International Symposium on Computational Geometry, pp. 1–15 (2018)
11. Bonnet, É., Miltzow, T.: Parameterized hardness of art gallery problems. In: 24th Annual European Symposium on Algorithms, pp. 1–17 (2016)
12. Bonnet, É., Miltzow, T.: An approximation algorithm for the art gallery problem. In: 33rd International Symposium on Computational Geometry, pp. 1–15 (2017)
13. Çağırıcı, O., Ghosh, S., Hliněný, P., Roy, B.: On conflict-free chromatic guarding of simple polygons. arXiv:1904.08624, June 2019
14. Chazelle, B.: Approximation and decomposition of shapes. In: Schwartz, J.T. (ed.) Advances in Robotics 1: Algorithmic and Geometric Aspects of Robotics, pp. 145–185 (1987)
15. Chvátal, V.: A combinatorial theorem in plane geometry. J. Comb. Theory **18**, 39–41 (1975)
16. Efrat, A., Har-Peled, S.: Guarding galleries and terrains. Inf. Process. Lett. **100**, 238–245 (2006)
17. Eidenbenz, S., Stamm, C., Widmayer, P.: Inapproximability results for guarding polygons and terrains. Algorithmica **31**, 79–113 (2001)

18. Erickson, L.H., LaValle, S.M.: An art gallery approach to ensuring that landmarks are distinguishable. In: Robotics: Science and Systems VII (2011)
19. Ghosh, S.K.: Approximation algorithms for art gallery problems in polygons. Discrete Appl. Math. **158**, 718–722 (2010)
20. Guibas, L.J., Motwani, R., Raghavan, P.: The robot localization problem. SIAM J. Comput. **26**, 1120–1138 (1997)
21. Guy, R.K.: Unsolved Problems in Number Theory. Springer, New York (1994). https://doi.org/10.1007/978-1-4899-3585-4
22. Hoffmann, F., Kriegel, K., Suri, S., Verbeek, K., Willert, M.: Tight bounds for conflict-free chromatic guarding of orthogonal art galleries. Comput. Geom. **73**, 24–34 (2018)
23. Katz, M.J.: A PTAS for vertex guarding weakly-visible polygons - an extended abstract. CoRR abs/1803.02160 (2018)
24. Katz, M.J., Morgenstern, G.: Guarding orthogonal art galleries with sliding cameras. Int. J. Comput. Geom. Appl. **21**, 241–250 (2011)
25. King, J., Kirkpatrick, D.G.: Improved approximation for guarding simple galleries from the perimeter. Discrete Comput. Geom. **46**, 252–269 (2011)
26. Kratochvíl, J., Křivánek, M.: On the computational complexity of codes in graphs. In: Chytil, M.P., Koubek, V., Janiga, L. (eds.) MFCS 1988. LNCS, vol. 324, pp. 396–404. Springer, Heidelberg (1988). https://doi.org/10.1007/BFb0017162
27. Lee, D.T., Lin, A.K.: Computational complexity of art gallery problems. IEEE Trans. Inf. Theory **32**, 276–282 (1986)
28. O'Rourke, J.: Art Gallery Theorems and Algorithms. Oxford University Press, Oxford (1987)
29. Smorodinsky, S.: Conflict-free coloring and its applications. In: Bárány, I., Böröczky, K.J., Tóth, G.F., Pach, J. (eds.) Geometry — Intuitive, Discrete, and Convex. BSMS, vol. 24, pp. 331–389. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-41498-5_12
30. Suri, S.: On some link distance problems in a simple polygon. IEEE J. Robot. Autom. **6**, 108–113 (1990)